

09/ 894 200

CoFC



PATENT
Attorney Docket No. 68970

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE


Patent No. 7,016,305

Confirmation No. 7077

Issued: March 21, 2006

Name of Patentee: Michelson et al.

Patent Title: METHOD AND APPARATUS
FOR DISTRIBUTING INFORMATION
WITHIN A PACKET SWITCHING SYSTEM



**REQUEST FOR CERTIFICATE OF CORRECTION OF
PATENT FOR PATENT OFFICE MISTAKE (37 C.F.R. § 1.322)**

Attn: Certificate of Correction Branch
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**Certificate
of Correction**
SEP 08 2006

Dear Sir:

It is requested that a Certificate of Correction be issued to correct Office mistakes found the above-identified patent. Attached hereto is a Certificate of Correction which indicates the requested correction. For your convenience, also attached are copies of selected pages (a) from the issued patent with errors highlighted, and (b) from the original application as filed June 27, 2001 and (c) Amendment A filed July 14, 2005, with the correct text/instructions.

SEP 08 2006

It is believed that there is no charge for this request because applicant or applicants were not responsible for such error, as will be apparent upon a comparison of the issued patent with the application as filed or amended. However, the Assistant Commissioner is hereby authorized to charge any fee that may be required to Deposit Account No. 501430.

Date: August 29, 2006

By

Kirk D. Williams, Reg. No. 42,229
One of the Attorneys for Applicants
CUSTOMER NUMBER 26327
The Law Office of Kirk D. Williams
1234 S. OGDEN ST., Denver, CO 80210
303-282-0151 (telephone), 303-778-0748 (facsimile)

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 7,106,305
DATED : March 21, 2006
INVENTOR(S) : Michelson et al.

It is certified that error(s) appear in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Col. 9, line 57, replace "41 1A-D" with -- 411 A-D --

Col. 10, line 64, replace "4101-L" with -- 410I-L --

Col. 10, line 66, replace "4101-L" with -- 410I-L --

Col. 11, line 17, replace "41D" with -- 411D --

Col. 13, line 40, replace "interfaces:" with -- interfaces; --

Col. 14, line 36, replace "repeatedly;" with -- repeatedly: --

Col. 14, line 46, replace "stare" with -- stage --

Col. 14, line 57, replace "claim 5" with -- claim 6 --

MAILING ADDRESS OF SENDER:

Kirk D. Williams, Reg. No. 42,229
Customer No. 26327
The Law Office of Kirk D. Williams
1234 S. Ogden Street, Denver, CO 80210

PATENT NO. 7,106,305

No. of additional copies

⇒ NONE (0)

08 08

puter-executable instructions to be executed by control logic 371 and/or data which is manipulated by control logic 371 for implementing functionality in accordance with certain embodiments of the invention. Storage devices 370 are another type of computer-readable medium, and typically comprise disk drives, diskettes, networked services, tape drives, and other storage devices. Storage devices 370 typically store computer-executable instructions to be executed by control logic 371 and/or data which is manipulated by control logic 371 for implementing an information distribution mechanism in accordance with certain embodiments of the invention.

Each SE-3 360 receives packets 361 and exchanges control messages 362 over one or more links with one or more SE-2 elements (not shown) such as SE-2 330 (FIG. 3B) via SE-2 interfaces 365. In other embodiments, data packets and control messages are transmitted over a common link or links, and/or communication interfaces have a folded topology. Additionally, SE-3 360 sends packets 388 and exchanges control messages 389 over one or more links with one or more output interface elements (not shown) such as Input/Output interface 390 (FIG. 2C) via I/O interfaces 385. Control logic 371 receives control packets containing information, and updates its data structure stored in memory 372. In one embodiment, there is an output queue 380 for each destination, or for each class of service for each destination. In one embodiment, each output queue 380 implements an information distribution mechanism according to the invention. In one embodiment, each input and/or output interface 365, 385 implements an information distribution mechanism according to the invention.

FIGS. 4A-D depict one embodiment of a packet switching system 400 implementing an information distribution mechanism according to the invention. The architecture of this illustrated switching fabric 400 is first described, which is then followed by a description of the accumulation and distribution of information. Although certain particular references herein are made to "flow control" information, the teachings and embodiments of the invention include all types of control and data information. Moreover, the invention is not limited to the particular switching fabric 400. Rather, the invention is extensible and FIGS. 4A-D illustrate the operation of one of an unlimited number of embodiments. The teachings of the invention and this particular embodiment may be adapted for other switching implementations, architectures, and computer and communications systems.

FIGS. 4A-D illustrate a packet switching fabric having multiple I/O interfaces 410A-P. Note, the illustration of the topology of packet switching fabric 400 represents the same I/O interfaces 410A-P on both the left and right sides of the figures for simplicity of illustration and ease of understanding, rather than drawing links back to one representation of I/O interfaces 410A-P. Also, switch elements SE-1 411A-D and SE-3 413A-D are illustrated separately; however in certain embodiments such as that illustrated in FIG. 1C, some of these switching elements 411A-D and 413A-D may be implemented in the same component. Moreover, other embodiments employ a different packet switch topology, such as, but not limited to a non-folded network, which provides some mechanism to convey flow control and other information from the output or egress portion of the packet switching back to the input or ingress portion.

In general, first stage switching elements 411A-D include a data structure (DS) 422A-D for storing received flow control information. For example, using any technique, flow control information is received from sources internal to

packet switching system 400, including from I/O interfaces 410A-P and/or their connected devices (not shown). The transmission of this flow control information to the traffic sources and/or other consumers is described in relation to FIGS. 4A-D, which show one embodiment of a set of paths over which flow control information is forwarded to I/O interfaces 410A-P.

The paths or routes that flow control information is distributed in FIGS. 4A-D are exemplary, and not limiting. Moreover, FIGS. 4A-D should be considered in conjunction with each other as the illustrated paths over which flow control information is transmitted according to the invention are selected to eliminate redundant flow control information arriving at the traffic sources or other destinations, and to minimize or eliminate the use of the same path by other switching elements, except, of course, where there are less paths to a next destination than flow control information sources (e.g., between third stage switching elements 413A-D and I/O interfaces 410A-P).

In one embodiment shown in FIG. 4A, first stage switching element 411A receives flow control information from connected I/O interfaces 410A-D and stores this information in data structure 422A. The source of the flow control information received from I/O interfaces 410A-D may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411A extracts a portion of the flow control data structure 411A, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412A. In one embodiment, first stage switching element 411A transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2 412A in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412A must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

Similarly, in one embodiment shown in FIG. 4B, first stage switching element 411B receives flow control information from connected I/O interfaces 410E-H and stores this information in data structure 422B. The source of the flow control information received from I/O interfaces 410E-H may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411B extracts a portion of the flow control data structure 411B, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412B. In one embodiment, first stage switching element 411B transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2 412B in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412B must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

Similarly, in one embodiment shown in FIG. 4C, first stage switching element 411C receives flow control information from connected I/O interfaces 410I-L and stores this information in data structure 422C. The source of the flow control information received from I/O interfaces 410I-L may be from any source internal to or external to packet

411 A-D

410I-L

410I-L

6444

08

11

switching system 400. Periodically, first stage switching element 411C extracts a portion of the flow control data structure 411C, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412C. In one embodiment, first stage switching element 411C transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2 412C in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412C must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

Similarly, in one embodiment shown in FIG. 4D, first stage switching element 411D receives flow control information from connected I/O interfaces 410M-P and stores this information in data structure 422D. The source of the flow control information received from I/O interfaces 410M-P may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411D extracts a portion of the flow control data structure 411D, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412D. In one embodiment, first stage switching element 411D transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2 412D in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412D must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

FIG. 5A illustrates an exemplary data structure 500, which, in one embodiment, corresponds to data structures 422A-D (FIGS. 4A-D). Data structure 500 is comprised of rows of I/O interfaces 501 and columns of groups of destinations 502, wherein a destination group corresponds to all the I/O interfaces connected to a particular first and third stage switching component. In one embodiment, a simple binary flag is used to indicate an XON or XOFF condition, while in other embodiments, an index, such as a numerical value, is used to indicate a level of traffic.

Different forms of data structures are used in different embodiments, including, but not limited to linked lists, trees, etc. Also, information is stored in data structures used in different embodiments, including for example, but not limited to source and destination pairings of control information. In one embodiment, a simple binary flag is used to indicate an XON or XOFF condition between as source-destination pair, while in other embodiments, an index, such as a numerical value, is used to indicate a level of traffic between a source-destination pair.

FIG. 5B illustrates an exemplary embodiment of a packet 550 used to distribute flow control information. As shown, packet 550 includes a header field 551, control field 552, a flow control address 553 used to indicate a context of the flow control information (e.g., to what source-destination pair or pairs the information belongs, or a position within a flow control data structure such as data structure 500 illustrated in FIG. 5A), flow control data field 554, and possibly other control and data fields 559. In one embodiment, only flow control information is included in packet 550. In one embodiment, flow control information is piggybacked in

12

packets 550 which contain other information. In one embodiment, the flow control address field 553 only contains a portion of the address of the information received in flow control data field 554, wherein the remainder of the address may be determined based on the path or from which element packet 550 was received. In one embodiment, flow control address field 553 indicates a start of sequence or sequences of information, and the actual flow address may be inferred by the receipt number in a predetermined sequence of received packets 550. In one embodiment, other control and data fields 559 contains error detection and/or error correction information.

FIG. 6A illustrates a process of one embodiment used by a component of one embodiment to update its flow control data structure. This component of the packet switching system is typically determined based on the topology of the packet switching system. For example, in a non-folded Benes network, this component is typically a first stage switching element. Processing begins at process block 600, and proceeds to process block 602, wherein flow control information is received (e.g., via signals, messages, packets, etc.). In process block 604, this flow control information is used to update (e.g., replace or modify current data) the flow control data structure, and processing returns to process block 602.

FIG. 6B illustrates a process of one embodiment used to forward flow control information collected in its data structure. Processing begins at process block 620, and proceeds to process block 622, wherein the process is initialized, which may include initializing a current position within the data structure, etc. Next, in process block 624, a packet to forward is received or created, which typically depends on whether the embodiment piggybacks flow control information or send flow control information in created packets. Next, as determined in process block 630, if it is time to send a blank flow control information (e.g., so as to not send flow control information at a maximum rate so as to allow for variations in bandwidth and timing between components), then in process block 632, a blank flow control indication is inserted in the packet. Otherwise, in process block 634, the flow control information at the current position within the flow control data structure is inserted into the packets (e.g., into flow control address and data fields 553 and 554 of FIG. 5B), and in process block 636, the current position is advanced. Next in process block 640, the packet is forwarded to a next component, typically to all components for simplicity of implementation and handling of error conditions, or possibly only to one or more designated next components.

FIG. 7A illustrates a process used by an intermediate stage switching element in one embodiment. Processing begins at process block 700, and proceeds to process block 702, wherein a packet with flow control information is received. Next, as determined in process block 704, if the flow control information came from a designated source, then in process block 706 the flow control information is extracted, and in process block 708, the flow control information is inserted in packets going to all next stage switching elements (or a subset thereof in one embodiment). Processing returns to process block 702.

FIG. 7B illustrates a process used in one embodiment for handling error conditions and/or reconfiguration commands. Processing begins at process block 720. If, as determined in process block 722 and error condition is detected that affects the current configuration of the flow control distribution, then in process block 724, the designated source of flow control information is updated. If, as determined in process

13

block 726, that a route update command has been received, then in process block 728, the designated source is updated per the received command in process block 728. Processing returns to process block 722.

FIG. 8A illustrates a process for forwarding flow control information of a final stage switching element of one embodiment. Processing begins at process block 800, and proceeds to process block 802, wherein a packet with flow control information is received. Next, in process block 804, this flow control information is extracted, and in process block 806, this flow control information is forwarded to each of the I/O interfaces and/or line cards. Processing returns to process block 802.

FIG. 8B illustrates a process for forwarding flow control information to a component which maintains a flow control data structure. Processing begins in process block 820, and proceeds to process block 822, wherein a flow control condition is identified (e.g., via a signal, packet, or other detection or communications mechanism). In process block 824, this flow control information is then sent to the first stage switching element (or other component) in one embodiment (and handled in one embodiment per the process illustrated in FIG. 6A). Processing returns to process block 822.

In view of the many possible embodiments to which the principles of our invention may be applied, it will be appreciated that the embodiments and aspects thereof described herein with respect to the drawings/figures are only illustrative and should not be taken as limiting the scope of the invention. For example and as would be apparent to one skilled in the art, many of the process block operations can be re-ordered to be performed before, after, or substantially concurrent with other operations. Also, many different forms of data structures could be used in various embodiments. The invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

What is claimed is:

1. A packet switch, comprising:
 - a plurality of input/output interfaces;
 - a benes switching fabric including a plurality of first stage switching elements, a plurality of second stage switching elements, and a plurality of third stage switching elements communicatively coupled via a plurality of paths according to a benes topology;
 - wherein each first stage switching element of the plurality of first stage switching elements includes a flow control storage mechanism for storing received flow control information extracted from information received from the corresponding input/output interfaces of said input/output interfaces to which said first stage switching element is coupled, and control logic for receiving said information, extracting said received flow control information from said information, and for updating the storage mechanism with said received flow control information;
 - wherein each of the first stage switching elements is configured to repeatedly sequence through said flow control information currently stored in the storage mechanism and to send a portion of said flow control information stored at a current location within said flow control information over one of said paths to a predetermined one of the plurality of second stage switching elements, such that each of the first stage switching elements is configured to send said flow control information to a different one of the plurality of second stage switching elements;

14

wherein each of the plurality of second stage switching elements is configured to receive and forward said portions of flow control information to each of the plurality of third stage switching elements; and

wherein each of the plurality of third stage switching elements are configured to send said portions of flow control information to each of the plurality of input/output interfaces for performing flow control operations in response to said flow control information.

2. The packet switch of claim 1, where the flow control storage mechanism includes a flow control data structure indexed by said input/output interfaces.

3. The packet switch of claim 2, where each said information includes an address within the flow control data structure at which to store said extracted flow control information.

4. The packet switch of claim 1, wherein the packet switch is configured to re-route said sending and forwarding of said flow control information in the benes switching fabric in response to an identified error within the benes switching fabric.

5. A packet switch, comprising:

a plurality of input/output interfaces;

a benes switching fabric including a plurality of first stage switching elements, a plurality of second stage switching elements, and a plurality of third stage switching elements communicatively coupled via a plurality of paths according to a benes topology;

wherein each first stage switching element of the plurality of first stage switching elements includes: means for maintaining a flow control data structure and for updating the flow control data structure with received flow control information from the corresponding input/output interfaces of said input/output interfaces to which said first stage switching element is coupled; and means for repeatedly retrieving a portion of said flow control information currently stored in the flow control data structure and sending the portion of said flow control information over one of said paths to a predetermined one of the plurality of second stage switching elements, such that each of the first stage switching elements is configured to send said flow control information to a different one of the plurality of second stage switching elements;

wherein each of the plurality of second stage switching elements includes means receiving and forwarding said portions of flow control information to each of the plurality of third stage switching elements; and

wherein each of the plurality of third stage switching elements includes means sending said portions of flow control information to each of the plurality of input/output interfaces for performing flow control operations in response to said flow control information.

6. The packet switch of claim 5, where the flow control data structure indexed by said input/output interfaces.

7. The packet switch of claim 5, where said received flow control information is accompanied by an address within the flow control data structure at which to store said flow control information.

8. The packet switch of claim 5, wherein the packet switch is configured to re-route said sending and forwarding of said flow control information in the benes switching fabric in response to an identified error within the benes switching fabric.

* * * * *

SEP 08

From Application filed 6-27-2007

have a folded topology. Additionally, SE-3 360 sends packets 388 and exchanges control messages 389 over one or more links with one or more output interface elements (not shown) such as Input/Output interface 390 (FIG. 2C) via I/O interfaces 385. Control logic 371 receives control packets containing information, and updates its data structure stored in memory 372. In one embodiment, there is an output queue 380 for each destination, or for each class of service for each destination. In one embodiment, each output queue 380 implements an information distribution mechanism according to the invention. In one embodiment, each input and/or output interface 365, 385 implements an information distribution mechanism according to the invention.

FIGs. 4A-D depict one embodiment of a packet switching system 400 implementing an information distribution mechanism according to the invention. The architecture of this illustrated switching fabric 400 is first described, which is then followed by a description of the accumulation and distribution of information. Although certain particular references herein are made to "flow control" information, the teachings and embodiments of the invention include all types of control and data information. Moreover, the invention is not limited to the particular switching fabric 400. Rather, the invention is extensible and FIGs. 4A-D illustrate the operation of one of an unlimited number of embodiments. The teachings of the invention and this particular embodiment may be adapted for other switching implementations, architectures, and computer and communications systems.

FIGs. 4A-D illustrate a packet switching fabric having multiple I/O interfaces 410A-P. Note, the illustration of the topology of packet switching fabric 400 represents the same I/O interfaces 410A-P on both the left and right sides of the figures for simplicity of illustration and ease of understanding, rather than drawing links back to one representation of I/O interfaces 410A-P. Also, switch elements SE-1 411A-D and SE-3 413A-D are illustrated separately; however in certain embodiments such as that illustrated in FIG. 1C, some of these switching elements 411A-D and 413A-D may be implemented in the same component. Moreover, other embodiments employ a different packet switch

Col. 9, line 57

From Application filed 6-27-2001

example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412A must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

5 Similarly, in one embodiment shown in FIG. 4B, first stage switching element 411B receives flow control information from connected I/O interfaces 410E-H and stores this information in data structure 422B. The source of the flow control information received from I/O interfaces 410E-H may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411B extracts a
10 portion of the flow control data structure 411B, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412B. In one embodiment, first stage switching element 411B transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2
15 412B in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412B must broadcast the flow control information to all third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

20 Similarly, in one embodiment shown in FIG. 4C, first stage switching element 411C receives flow control information from connected I/O interfaces 410I-L and stores this information in data structure 422C. The source of the flow control information received from I/O interfaces 410I-L may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411C extracts a
25 portion of the flow control data structure 411C, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412C. In one embodiment, first stage switching element 411C transmits the flow control information to all second stage

Col. 10, line 64

Col. 10, line 66

From Application filed 6-27-2001

switching elements 412A-D, and only the designated second switching element (SE-2 412C in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412C must broadcast the flow control information to all
5 third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

Col. 11, line 17 Similarly, in one embodiment shown in FIG. 4D, first stage switching element 411D receives flow control information from connected I/O interfaces 410M-P and stores this information in data structure 422D. The source of the flow control information
10 received from I/O interfaces 410M-P may be from any source internal to or external to packet switching system 400. Periodically, first stage switching element 411D extracts a portion of the flow control data structure 411D, and sends it over a path to a predetermined or programmed one of the second stage switching elements 412A-D, in this example to second stage switching element 412D. In one embodiment, first stage
15 switching element 411D transmits the flow control information to all second stage switching elements 412A-D, and only the designated second switching element (SE-2 412D in this particular example), forwards the information to a next stage. As flow control information is being distributed in this example to all I/O interfaces 410A-P, second stage switching element 412D must broadcast the flow control information to all
20 third stage switching elements 413A-D, which then forwards this information to all I/O interfaces 410A-P.

FIG. 5A illustrates an exemplary data structure 500, which, in one embodiment, corresponds to data structures 422A-D (FIGs. 4A-D). Data structure 500 is comprised of rows of I/O interfaces 501 and columns of groups of destinations 502, wherein a
25 destination group corresponds to all the I/O interfaces connected to a particular first and third stage switching component. In one embodiment, a simple binary flag is used to indicate an XON or XOFF condition, while in other embodiments, an index, such as a numerical value, is used to indicate a level of traffic.

From Amendment A filed 7-14-2005

In re MICHELSON ET AL., Application No. 09/894,200
Amendment A

Amendments to the Claims:

The listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

Claim 1 (currently amended): A packet ~~switching system~~ switch, comprising:
a plurality of input/output interfaces;
a benes switching fabric including a plurality of first stage switching elements, a
plurality of second stage switching elements, and a plurality of third stage switching elements
communicatively coupled via a plurality of paths according to a benes topology;

a plurality of first components, wherein each first stage switching element of the
plurality of first stage switching elements components including: includes a flow control
storage device mechanism for storing received flow control information extracted from
information received from the corresponding input/output interfaces of said input/output
interfaces to which said first stage switching element is coupled, and control logic for
receiving said information, extracting said received flow control information from said
information, and for updating the storage mechanism with said device with indications of the
received flow control information;

a plurality of second components;

a plurality of paths between each of the plurality of first components and each of the
plurality of second components;

wherein each of the first stage switching elements components is configured to
repeatedly sequence through a portion of the storage device, said flow control information
currently stored in the storage mechanism and to send a portion of said flow control
information stored at a current location within said flow control information the portion of the
storage device over the plurality of paths over one of said paths to each a predetermined one
of the plurality of second stage switching elements, such that each of the first stage switching
elements is configured to said send said flow control information to a different one of the
plurality of second stage switching elements; components; and

From Amendment A filed 7-14-2005

In re MICHELSON ET AL., Application No. 09/894,200
Amendment A

Claim 31 (new): A packet switch, comprising:

a plurality of input/output interfaces;

a benes switching fabric including a plurality of first stage switching elements, a plurality of second stage switching elements, and a plurality of third stage switching elements communicatively coupled via a plurality of paths according to a benes topology;

wherein each first stage switching element of the plurality of first stage switching elements includes: means for maintaining a flow control data structure and for updating the flow control data structure with received flow control information from the corresponding input/output interfaces of said input/output interfaces to which said first stage switching element is coupled; and means for repeatedly: retrieving a portion of said flow control information currently stored in the flow control data structure and sending the portion of said flow control information over one of said paths to a predetermined one of the plurality of second stage switching elements, such that each of the first stage switching elements is configured to said send said flow control information to a different one of the plurality of second stage switching elements;

wherein each of the plurality of second stage switching elements includes means receiving and forwarding said portions of flow control information to each of the plurality of third stage switching elements; and

wherein each of the plurality of third stage switching elements includes means sending said portions of flow control information to each of the plurality of input/output interfaces for performing flow control operations in response to said flow control information.

Claim 32 (new): The packet switch of claim 31, where the flow control data structure indexed by said input/output interfaces.

Claim 33 (new): The packet switch of claim 32, where said received flow control information is accompanied by an address within the flow control data structure at which to store said flow control information.